AD-A123 359 BIBLIOGRAPHY(U) ILLINOIS UNIV AT URBANA COORDINATED SCIENCE LAB DEC 82 N00014-75-C-0612 1/1

UNCLASSIFIED F/G 5/2 NL

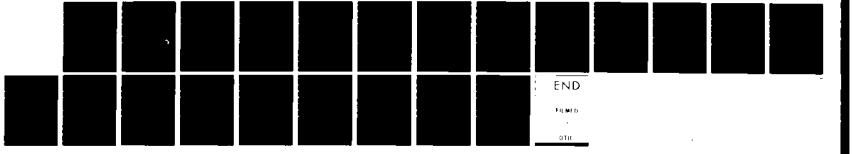

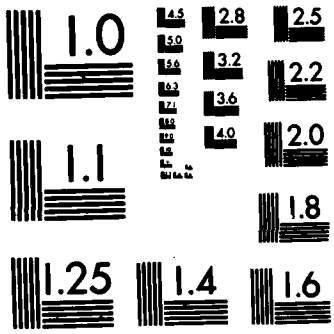

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

WORKING PAPER 0

Bibliography

Advanced Automation Group
Coordinated Science Laboratory
University of Illinois at Urbana-Champaign
Urbana, IL 61801

December 1982

Abstract

This bibliography lists working papers and reports written by members of the Advanced Automation Group of the Coordinated Science Laboratory of the University of Illinois.

This work is supported by the Office of Naval Research under Contract N00014-75-C-0612.

This bibliography lists working papers and reports written by members of the Advanced Automation Research Group of the Coordinated Science Laboratory of the University of Illinois. Requests for orders can be sent to:

Publications
Advanced Automation Research Group
Coordinated Science Laboratory
University of Illinois
Urbana, IL 61801

Some of these reports are also available from other sources, including conference reports and the National Technical Information System.

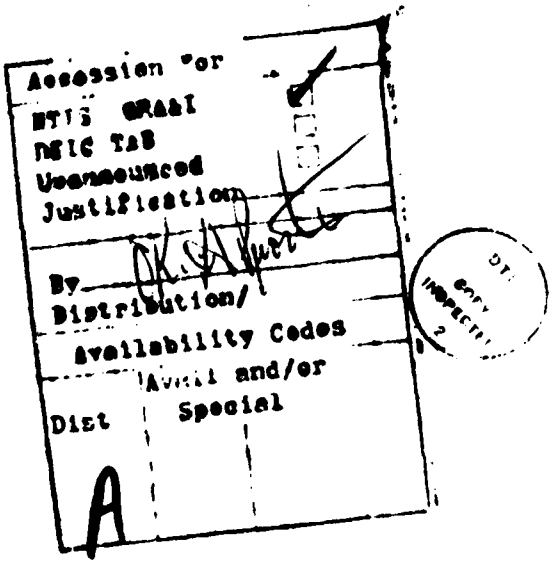

WORKING PAPERS

WP-0 Bibliography
anonymous; January 1981
20 pages; $0.00 (no charge)

This bibliography lists working papers and reports written by members of the Advanced Automation Group of the Coordinated Science Laboratory of the University of Illinois.

WP-1 The LISP Editor
R. P. Gabriel and Tim Finin; February 1975
6 pages; $0.70 (on line: edit.doc [5,720,document])

This report describes a LISP editor which allows a user to write, edit, test, and file functions without leaving the LISP environment. The editor also has powerful matching abilities, undoable commands, and a sophisticated file manager.

WP-2 PSEUDOCONNIVER: Some CONNIVER Demons for Those Who Don't Need Anything Else
R. P. Gabriel and Steve Weissman; February 1975
5 pages; $0.70 (on line: conv.doc [5,720,document])

This memo describes a package of LISP functions which allow the user to have some of the data base manipulating facilities found in CONNIVER without having the associated frills and overhead.

WP-3 Navy Data Base Read Package
Paul Rutter; 1975
9 pages; $0.80

The purpose of this report is to explain the structure of the Navy aircraft maintenance records data base and to document some functions which have been written to do low level accessing and searching of this data base.

WP-4 The CSL Vision Library
V. Jones and W. E. Snyder; March 1975
23 pages; $1.40

This report describes routines and programs available to facilitate using the CSL Vision System. The functions are written in either BLISS or BLISS compatible MACRO. The programs are already compiled and can be executed via the monitor RUN command.

WP-5 Real World Data on the 3-M Data Base
Forrest Conrad; January 1976
11 pages; $0.90

Based on a report of actual usage of the 3-M Data Base for Aircraft, this document provides information on: (1) frequently investigated categories of concern, (2) corresponding types of sentences the PLANES system must be able to answer, and (3) a summary of the most frequently used abbreviations.

WP-6 TECOX Macros
Tim Finin; April 1976
7 pages; $0.70 (on line: xmacro.doc [1000,130,teco])

This memo describes a set of TECO macros which provide a variety of helpful editing functions. A number of the macros are specially designed for TECOing LISP code, enabling one to deal with "s-expressions" as units. Other macros are more generally useful for doing such things as manipulating chunks of text, simplifying I/O, and displaying portions of the buffer.

WP-7 Optimizing Procedures: Algorithmic Improvement and Heuristic Generation
Paul Rutter; October 1976
48 pages; $2.40

Two techniques for improving time-bounded performance of a given procedure are discussed: Algorithmic Improvement and Heuristic Generation. Examples of both techniques are analyzed in the successive optimization of a particular procedure. A programming system is proposed which will be used to investigate what knowledge and skills are necessary to carry out high level optimization automatically.

WP-8 Human Factor Considerations For Interactive Dialogs Using Non-exotic Terminals
Harry Tennant; March 1977
14 pages; $1.00

This paper is a collection of miscellaneous features that can be implemented fairly easily on non-exotic alphanumeric and garden variety video terminals. The purpose of these features is to facilitate the use of the computer system without additional hardware investment. The software will depend upon the habitablity of the online system at the start.

WP-9 Source-to-Source STATIC Transformations for LISP Programs
Paul Rutter; March 1977
36 pages; $1.90

A list of Source-to-Source transformations is given which improve the speed of LISP programs. The transformations are written in the pattern matching language used by the author in a program manipulation system written for his thesis. Only the static transformations from the thesis are listed here (those which can be applied knowing only the program text). Included are transformations which improve arithmetic and conditional expressions, do-loops, usage of LISP functions, perform replacement of function calls by their bodies, and carry out recursive to iterative transformation.

WP-10 Design and Construction of a Natural Language Front End for a Large Relational Data Base
D. L. Waltz and Bradley A. Goodman; October 1977
31 pages; $1.70

We present a model for processing English requests for information from a relational data base. The model has as its main steps (a) locating semantic constituents of a request; (b) matching these constituents against larger template called concept case frames; (c) filling in the concept case frame using information from the user's request, from the dialogue context and from the user's responses to questions posed by the system; and (d) generating a formal data base query using the collected information. A method is suggested for constructing the components of such a natural language processing system for an arbitrary relational data base. The model has been applied to a large data base of aircraft flight and maintenance data to generate a system called PLANES; examples are drawn from this system.

WP-11 The Evaluation of Natural Language Question Answers
Harry Tennant; May 1979
29 pages; $1.60

The current practice for presenting research in natural language processing is to describe the techniques used and include a list of successfully analyzed queries. While it is difficult to conduct a thorough and controlled evaluation, we feel that the lack of information on system performance has had an adverse effect on the development of the field. This paper examines the problems of evaluation and description of natural-language processors. It makes suggestions on which parameters should be measured, which should be controlled, and how evaluation experiments should be designed.

WP-12 A Computational Look at Point of View in the Comprehension of Narrative
Jeff Gibbons; May 1978
21 pages; $1.30

In this paper we propose to apply the methodology of the computational paradigm to the literary notion of point of view as it relates to the process of comprehending narrative text. The computational paradigm comes from the field of artificial intelligence [Winston,1977; Bobrow and Collins,1975] and is discussed in section II. Section III develops a model of point of view within the concept of narrative and section IV describes its utility. Further research questions are posed in section V, in this section we also state our understanding of the narrative concept.

WP-13 Browsing and Alerting in Large Data Bases
Douglas D. Dankel II; February 1978
14 pages; $1.00

Superceded by working paper 17.

The formation of data bases containing information on mechanical systems for trouble shooting purposes has become increasingly popular and important. Examination of these data bases by humans can be very costly. A system called BROWSER is currently under development at the University of Illinois at Urbana-Champaign to heuristically search a data base containing information on Navy aircraft with little or no human intervention. BROWSER searches the data base guided by models and heuristics looking for interesting patterns or configurations. The user is then notified of the existence of these patterns and if the patterns are of interest, BROWSER can then create alerters, which monitor new data looking for the occurrence of the patterns.

WP-14 The Planes Database
Harry Tennant; July 1978
86 pages; $3.90

The purpose of this paper is to collect the information about the structure and content of the database for Planes into one document. The relevant information that is not included in this paper is referenced. An individual desiring to become familiar with the Planes database should start with this document. It may also serve as a reference for those who are familiar with the database.

WP-15 Augmenting ATNs
Tim Finin and George Hadden; August 1977
10 pages; $0.90

A short version of this paper appears in the Proceedings of the International Joint Conference on Artificial Intelligence-77.

In developing the PLANES natural language system we encountered several drawbacks and shortcomings in the standard Augmented Transition Network (ATN) model. Consequently, we have augmented and extended the model to alleviate some of these problems. We have also developd an optimizing compiler which translates ATN networks into LISP functions and an interactive program which facilitates extending and modifying an ATN grammar.

WP-17 Browsing in Large Data Bases
Douglas D. Dankel II; February 1979
24 pages; $1.40

Key words and phrases: heuristic search, pattern identification, data modeling, data bases, pattern recognition, alerters

A short version of this paper appears in the Proceedings of the International Joint Conference on Artificial Intelligence-79.

The formation of data bases containing information on mechanical systems for trouble shooting purposes has become increasingly popular and important. Examination of these data bases by humans can be very costly. A system called BROWSER is currently under development to heuristically search a data base containing information on Navy aircraft with little or no human intervention. BROWSER searches the data base guided by models and heuristics looking for interesting patterns or configurations. The user is then notified of the existence of these patterns.

WP-18 Experience with the Evaluation of Natural Language Question Answerers
Harry Tennant; February 1979
21 pages; $1.30

Key words and phrases: natural language processing, question answerers, evaluation, completeness, coverage

A short version of this paper appears in the Proceedings of the International Joint Conference on Artificial Intelligence-79.

Research in natural language processing could be facilitated by thorough and critical evaluations of natural language systems. Two measurements conceptual and linguistic completeness are defined and discussed in this paper. Testing done on two natural language question answerers demonstrated that the conceptual coverage of such systems should be extended to better satisfy the needs and expectations of users. Three heuristics are presented that describe how conceptual coverage of question answerers should be extended.

WP-19 JETS: Achieving Completeness through Coverage and Closure
Tim Finin, Brad Goodman and Harry Tennant; February 1979
27 pages; $1.50

Key words and phrases: natural language understanding, relational data bases, very large data bases, natural language front ends, English dialogue, frames, coverage, completeness, closure, JETS, PLANES

A version of this paper appears in the Proceedings of the International Joint Conference on Artificial Intelligence-79.

Work in progress on JETS the successor to PLANES is described. JETS is a natural language question answering system that is intended to interface users to a large relational data base. The architecture is designed to extend the conceptual coverage of JETS to better meet the conversational and data base usage requirements of users. The implementation of JETS is designed to gain a high degree of closure over concept manipulation contributing to a solution to the problems of perspicuity and scale. Specific examples are given of concept manipulation through the implied relationships of modification and of an approach to problem-solving through the use of frames.

WP-20 Visual Analog Representation for Natural Language Understanding
David L. Waltz and Lois Boggess; February 1979
28 pages; $1.60

Key words and phrases: natural language understanding, language and perception, scene description, representation of knowledge, frames

A version of this paper appears in the Proceedings of the International Joint Conference on Artificial Intelligence-79.

In order for a natural language system to truly "know what it is talking about" it must have a connection to the real world correlates of language. For language describing physical objects and their relations in a scene a visual analog representation of the scene can provide a useful target structure to be shared by a language understanding system and a computer vision system. This paper discusses the generation of visual analog representations from input English sentences. It also describes the operation of a LISP which generates such a representation from simple English sentences describing a scene. A sequence of sentences can result in a fairly elaborate model. The program can then answer questions about relationships between the objects even though the relationships between the objects may not have been explicit in the original scene description. Results suggest that the direct testing of visual analog representations may be an important way to bypass long chains of reasoning and to thus avoid the combinatorial problems inherent in such reasoning methods.

WP-21 The Semantic Interpretation of Noun-Noun Modification
Timothy W. Finin; May 1979
42 pages; $2.10

Key words and phrases: natural language understanding, semantics, semantic interpretation, nominal compounds, nominalization, noun-noun modification, modification, linguistics, noun phrase

This paper is a proposal for a Ph.D. thesis concerned with the semantic interpretation of noun-noun modification. The study of this topic is being coordinated with the development of the JETS natural language question answering system. One of the goals of this research is a computer program which will interpret instances of noun-noun modification within the domain of discourse of JETS. Related work on noun-noun modification in the disciplines of Linguistics and Artificial Intelligence is described and contrasted to the proposed research. The basic approach being taken is described and the scope of the work is outlined.

WP-22 Problem Solving in a Natural Language Environment
Bradley A. Goodman; July 1979
54 pages; $2.60

Key words and phrases: natural language understanding, problem solving, vague and complex requests, frames, relational data bases, natural language front ends, JETS, PLANES

The kinds of requests that can be currently handled by natural language data base systems are constrained mainly to simple queries to retrieve information from the data base. The requests must be completely specified by the user (though certain information can be assumed from past context). This paper is a proposal for a Ph.D. thesis that explores requests of a more complicated nature. The goal is to take vague and complex requests from users and turn them into well-defined problems. Missing information will be filled in through world knowledge or from the current dialog context. The transformation for the request into a well- defined problem and the generation of a plan to solve the problem will be guided by a set of problem solving frames.

WP-23 Relating Images Concepts and Words
David L. Waltz; May 1979
29 pages; $1.60

A version of this paper appears in Proceedings of the NSF Workshop on the Representation of Three-Dimensional Objects, University of Pennsylvania May 1979.

Examination of verbal descriptions of objects suggests that we use hierarchical structures for shape description; the highest levels of the hierarchy provide a general object framework or breakdown into component parts and a description of each part by analogy to a well-understood set of shapes called prototypes. Lower levels of the hierarchy provide refinement of the analogies and ways in which shapes deviate from the prototypes. The set of prototypes on which the analogies are based contains many common objects especially natural objects and the parts of the human body plus certain shapes with special symmetry properties. It is argued that no single 3-D representation scheme is natural for all members of this set of prototypes and that since unfamiliar objects are described with respect to the basic set of shapes these objects will have varying shape representation schemes also.

WP-24 Generating and Understanding Scene Descriptions
David L. Waltz; March 1980
24 pages; $1.40

This paper appears in Joshi, Sag and Webber (eds.), Elements of Discourse Understanding, Cambridge University Press, 1980.

This paper explores design issues for a system which has both vision and language in particular a system which addresses both the problem of selecting appropriate words and sentences to describe a particular perceptual event and the related problem of making appropriate inferences about a natural language description of a perceptual event. It argues that perception is basically a description-building process and that the understanding of scene descriptions is ultimately based on our ability to first use scene descriptions to drive processes of "picture-building" and then to drive processes of "event-simulation" which animate the "pictures" we build to mimic the dynamics of the world.

WP-25 The Semantic Interpretation of Nominal Compounds
Timothy W. Finin; April 1980
8 pages; $0.80

A version of this paper appears in the Proceedings of the American Association for Artificial Intelligence, Stanford University, August 1980.

Key words and phrases: natural language, semantic interpretation, compounding, noun-noun modification, linguistics

This paper briefly introduces an approach to the problem of building semantic interpretations of nominal compounds i.e. sequences of two or more nouns related through modification. Examples of the kinds of nominal compounds dealt with are: "engine repairs" "aircraft flight arrival" "aluminum water pump" and "noun-noun modification".

WP-26 Syntactic Analysis in Jets
Harry Tennant; May 1980
12 pages; $0.90

Key words and phrases: natural language processing, syntax, ellipsis, anaphora

Our experience with PLANES has lent evidence to the opinion that syntactic rules ought to be included in a natural language processing system. They were not explicitly included in PLANES with desirable and undesirable results. The desirable result was that PLANES was quite capable at interpreting elliptical and non-grammatical utterances which occur fairly frequently in user interactions. The undesirable result was that PLANES could not interpret many grammatical utterances whose interpretation would have been greatly facilitated by using syntactic information. This report is the result of an investigation and experiments into using an extension of the LUNAR grammar for JETS the successor to PLANES. The goal is to gain the regularities of syntactic analysis but still allow for elliptical and "non-grammatical" utterances.

WP-27 The State-of-the-Art in Natural Language Understanding
David L. Waltz; December 1980
35 pages; $1.90

A version of this paper will appear in Strategies for Natural Language Processing, NJ, 1982.

Research in computer understanding of natural language has led to the construction of programs which can handle a number of different types of language including questions about the contents of data bases stories and news articles dialogues and scene descriptions. This research draws on and has in turn had an effect on many other research areas including software engineering linguistics psychology philosophy and knowledge representation. This paper provides a brief history and overview of the field along with examples and explanations of the operation of several natural language understanding programs. The limitations of our current technology are discussed and assessments are given of the most promising current research directions.

WP-28 Understanding Event Descriptions
David L. Waltz; May 1980
17 pages; $1.10

A similar paper appeared in the Proceedings of the 18th Annual Meeting of the Association for Computational Linguistics and Parasession on Topics in Interactive Discourse, June 1980.

Key words and phrases: natural language processing, common sense language and perception, plausibility judgment

How do we know that an event description makes sense? It is argued that a hearer must create and run "event simulations" to check the consistency and plausibility of a "picture" constructed from a speaker's description.

A speaker must also run similar event simulations on his own descriptions in order to be able to judge when the hearer has been given sufficient information to construct an appropriate "picture" and to be able to respond appropriately to the hearer's questions about or responses to the scene description.

In this paper I explore some simple scene description examples in which a hearer must make judgments that involve reasoning about scenes space common-sense physics cause-effect relationships etc. My primary concern is to flesh out our understanding of just what the mechanisms must accomplish: what information will be available to them and what information must be found or generated to account for the inferences we know are actually made. I also propose some mechanisms for dealing with scene and event descriptions.

WP-29 Toward a Detailed Model of Processing for Language Describing the Physical World
David L. Waltz; June 1982
6 pages; $0.70

This paper appears in the Proceedings of the International Joint Conference on Artificial Intelligence-81.

This paper explores the problem of judging whether or not an English sentence could correspond to a real world situation or event which is literally physically plausible and the related problem of representing the different possibly physical situations. The judgement of plausibility can be made at a high level by checking semantic marker restrictions on verb case frame constituents. Often however plausibility judgment can only be based on the results of an attempt to construct (imagine) a scene that corresponds to the sentence and which does not violate "common sense" (i.e. relevant physical laws and expected stereotyped behavior). Methods are represented for constructing representations for different scenes which could correspond to a sentence. These methods incorporate (1) "subscripts" (sequences of scenes which comprise an event with attached preconditions and postconditions) to express different verb senses; (2) object representations which express properties such as shape size weight strength and behavior under common conditions; (3) physical laws encoded as constraints on behavior; (4) representation of context; and (5) robot problem solving-like methods to fit all this material together.

WP-30 Generalizations Based on Explanations
Gerald DeJong; June 1982
3 pages; $0.60

This paper appears in the Proceedings of the International Joint Conference on Artificial Intelligence-81.

This paper describes a new project at the University of Illinois in

computer learning. The phenomenon under study is a kind of "insight learning" of procedural schemata. The system described here is designed to grasp some principle underlying a natural language input. The underlying principle results in a new schema for the system. Once acquired the schema serves the same purpose as the other schemata in the system: it aids in processing future natural language inputs.

The neutral term "schema" rather than "frame" (Minsky (1975) Charniak (1976)) or "script" (Schank & Abelson (1977)) is used to refer to knowledge chunks because a frame (which is used to describe static objects as well as progressions of world situations) is too general a notion and the notion behind a script is too specific. Scripts cannot be made to fit a novel situation nor are they designed to represent the more abstract concepts that this system will need.

The process that the system uses is called explanatory schema acquisition. The basic idea behind it is that the causal connections in an understood representation of a new input can be used to propose and propagate constraints on slot fillers. That is from one particular instance of a situation the system can "reason out" the general structure underlying that instance. The system is therefore capable of learning from just one example.

WP-31 An Activation/Inhibition Network Cell
Jordan Pollack; January 1982
15 pages; $2.00

Spreading Activation and Lateral Inhibition are highly parallel information processing techniques which have been used with some success in simulating human cognitive faculties on computers. Unfortunately, these simulations run like molasses on serial machines. This paper describes the design of a VLSI-based architecture for the parallel simulation of activation and inhibition.

WP-32 Artificial Intelligence
David L. Waltz; April 1982
39 pages; $2.00

A substantially revised version of this article appears in Scientific American 247, 4, October 1982, 118-133.

This article is intended to be a brief introduction to AI. It surveys the kinds of results AI can achieve by explaining the operation of a number of AI programs for natural language understanding, game playing, reasoning, vision, and learning. Most of the programs discussed have been widely referenced and are well-known to AI researchers.

WP-33 Event Shape Diagrams
David L. Waltz; June 1982
4 pages; $0.80

This paper also appears in Proc. National Conference on Artificial Intelligence, August 1982.

"Event shape diagrams" are proposed as a representation for capturing the nuances of meaning of verbs that describe similar events. These diagrams represent timing, causal relationships between case roles, and typical value ranges for role fillers. Event shape diagrams are expressed in terms of primitive predicates and timing information that we believe could be computed by perceptual systems, and are intended to be a step toward the eventual connection of language systems to perceptual (vision, hearing and touch) sensing systems. The diagrams are capable of representing modification of verbs by adverbs, can support judgments of the degree of plausibility of various interpretations of a sentence's meaning, and may be useful in figuring out the meaning of certain kinds of metaphors.

WP-34 Automatic Schema Acquistion in a Natural Language Environment
Gerald DeJong; June 1982
4 pages; $0.80

This paper also appears in Proc. National Conference on Artificial Intelligence, August 1982.

This paper outlines an approach to schema acquisition. The approach, called explanatory schema acquisition is applicable in problem solving situations and is heavily knowledge-based. Basically, learning is viewed as a fundamental part of the understanding process. Understanding a situation for which there is no existing schema involves generalizing the new event into a nascent schema. The new schema is then available to aid in future processing and can be further refined via that processing. This approach to learning is unique in several respects: it is not inductive and so is capable of one trial learning, it does not depend on failures to drive the learning process, and it is incremental and learns comparatively slowly. The learning procedure is outlined briefly with an example, a taxonomy of situations involving explanatory schema acquistion is given, and there is a brief discussion on the scope of the learning mechanism.

WP-35 Natural Language Processing Using Spreading Activation and Lateral Inhibition
Jordan Pollack and David L. Waltz; August 1982
13 pages; $2.00

A shorter version of this paper is included in the Proc. of the 1982 Cognitive Science Society Conference, August 1982.

The knowledge needed to process natural language comes from many sources. While the knowledge itself may be broken up modularly, into knowledge of syntax, semantics, etc., the actual processing should be completely integrated. This form of processing is not easily amenable to the type of processing done by serial "von Neumann" computers. This work in progress is an investigation of the use of a highly parallel, spreading activation and lateral inhibition network as a mechanism for integrated natural language processing.

WP-36 Understanding Novel Language
Gerald F. DeJong and David L. Waltz; July 1982
41 pages; $2.10

Abstract not available.

REPORTS

R-635 On the Implementation of an Information System for Semantic Modeling and Computer Understanding
Steven Jay Weissman; December 1973
46 pages; $2.30

R-641 Machine Perception of Objects with Curved Surfaces
Yean-Hsi Chang; March 1974
$0.50

R-642 A Study of Computer Aided Manufacturing
Tony C.H. Woo; March 1974
$0.50

R-765 Improving Programs by Source-to-Source Transformations
Paul E. Rutter; April 1977
114 pages; $5.00

Also available from DDC as ADA041776.

R-798 A Model for a Natural Language Data Base System
Bradley A. Goodman; October 1977
113 pages; $5.00

Also available from DDC as ADA057641.

R-805 On Computer Stereo Vision with Wire Frame Models
David J. Burr; December 1977
117 pages; $5.00

R-851 Visual Recognition of Artifacts by Computer
Charles J. Jacobus; August 1979
175 pages; $5.00

T-34 The PLANES System: Natural Language Access to a Large Data Base
D. L. Waltz, T. Finin, F. Green, F. Conrad, B. Goodman and G. Hadden;
November 1976
140 pages; $5.00

T-37 Getting the Gist: A Computational Theory of Sentence Understanding
Shigeki Nakajima; December 1976
38 pages; $2.00

T-38 BROWSER: A User Oriented Information Retrieval System
Forrest P. Conrad; December 1976
45 pages; $2.30

T-39 Implementation of a Query Language Based on the Relational Calculus
Fred R. Green Jr.; December 1976
56 pages; $2.70

T-48 An Interpreter and Compiler for Augmented Transition Networks
Tim Finin; July 1977
104 pages; $4.60

Also available from DDC as ADA047505.

T-49 NETEDI: An Augmented Transition Network Editor
George D. Hadden; July 1977
66 pages; 3.10

Also available from DDC as ADA047506.

T-75 Computational Interpretation of English Spacial Prepositions
Lois Carolyn Boggess; February 1979
158 pages; $5.00

It seems clear to anyone who pays attention to the use of prepositions in language that any one preposition when used to describe the spatial relationship between different objects can produce strikingly different mental models for different objects. The mental model produced by the description "a bowl on a table" seems to be somewhat different from that

produced by "a poster on a wall" which in turn is somewhat different from "a shelf on a wall" which again is different from "a fly on a ceiling".

It is the contention of this paper that the preposition in conjunction with a small set of features of the objects (mostly perceptual features) can account for such variations in spatial relations. The thesis discusses a means of taking English-language descriptions involving prepositions and their semantic subjects and objects and deriving a three-dimensional model of the spatial relationships of the subject and object.

The relationship of some of the spatial prepositions to a coordinate system is explored as well as canonical definitions for prepositions based on analyses of descriptions using "neutral" subjects and/or objects ("whatchamacallit" "you-know-what" and so on).

Examples are taken from a simple program which accompanies the theory. The program is supplied with approximate descriptions of the shapes of a variety of objects. Each preposition in the program has one definition (e.g. there is only one procedure for on rather than several--ON1 ON2 ON3 and so on); in general the definition is made up of several components each of which is responsive to a perceptual characteristic of the semantic subject or object.

The program takes extended descriptions involving many objects each of which is incorporated into the overall model. Once an object has been described it is possible to interrogate the model about the relation of that object to any other in the model without recourse to inference rules of the following kind: "if A is on B and B is in C then A is (probably) in C."

T-96 The Semantic Interpretation of Nominal Compounds
Timothy Wilking Finin; March 1980
161 pages; $5.00

This report deals with one aspect of enabling machines to communicate with people in a natural language. The particular problem which is the focus of this work is the interpretation of nominal compounds i.e. sequences of two or more nouns related through modification. Examples of the kinds of nominal compounds dealt with are: "engine repairs" "aircraft flight arrival" "aluminum water pump" and "noun-noun modification".

The interpretation of nominal compounds can be divided into three intertwined subproblems: lexical interpretation (mapping words into concepts) modifierN parsing (discovering the structure of strings with more than two nominals) and concept modification (assigning an interpretation to the modification of one concept by another). This last problem is the central one. The essential feature of this form of modification is that the underlying semantic relationship which exists between the two concepts is not explicit. Moreover a large number of

relationships might in principal exist between the two concepts. The selection of the most appropriate one depends on a host of semantic pragmatic and contextual factors.

As a part of this research a computer program was written to build an appropriate semantic interpretation when given a string of nouns. This program has been designed as one component of the natural language question answering system JETS. The interpretation is done by a set of semantic interpretation rules. Some of the rules are very specific capturing the meaning of idioms and canned-phrases. Other rules are very general representing fundamental case-like relationships which can hold between concepts. A strong attempt has been made to handle as much as possible with the more general highly productive interpretation rules.

The approach has been built around a frame-based representational system which represents concepts and the relationships between them. The concepts are organized into an abstraction hierarchy which supports inheritance of attributes. The same representational system is used to encode the semantic interpretation rules. An important part of the system is the concept matcher which given two concepts determines whether the first describes the second and if it does how well.

T-100 A Cellular Automata Approach to Computer Vision and Image Processing
George D. Hadden; September 1980
145 pages; $5.00

This dissertation describes a system called HEXVIS which performs operations on scenes. All operations are carried out in a hexagonal array of cellular automata-like objects which corresponds to that scene. The system can perform (for instance) the following tasks: recognition of edges corners and axes of symmetry texture discrimination determination of areas and generation of Voronoi tessellations.

First the scene is embedded in the hexagonal array then the cells pass messages describing the cells' contents to their neighbors which in turn pass them on. As these messages pass through cells they can interact with each other and with the contents of the cell in which they find themselves. The cells all perform the same operation or group of operations in parallel on their visiting messages. As a consequence of these operations the states of some cells change in a way which indicates that they correspond to "interesting" parts of the scene.

This process can be repeated recursively using the altered states of the cells as new messages to be broadcast.

In addition a new algorithm is presented which shrinks binary scenes in a hexagonal array. It is proved that all scenes with (at most) simply connected holes are transformed into a set of isolated points each corresponding to a connected object in the original scene. This shrinking algorithm is embedded in HEXVIS.

T-102 Browsing in Data Bases
Douglas D. Dankel II; November 1980
297 pages; $5.00

Consider the tasks of trying to find similarities between the causes of the crashes of several aircraft of a common type the causes or early warning signs for various diseases the characteristics of stocks with high growth potential or advanced warnings of severe weather conditions. All of these tasks require a large amount of data and valuable time spent shifting through the data. A browsing computer system could also perform these tasks.

This thesis examines one possible organization for a browsing system containing models and heuristics. The models describe the organization of the data base and the objects from which the data was gathered. The heuristics knowledge on important features which might exist within the data base and techniques for locating these features.

An implementation of the organization was performed using a data base describing maintenance performed on Navy aircraft. This computer system BROWSER attempts to find recurring time-sequences of maintenance and major differences in the types of maintenance performed on different groups of aircraft. The user is then notified of these discoveries in the hope that changes will be made to improve the performance of the aircraft.

T-103 Evaluation of Natural Language Processors
Harry R. Tennant; November 1980
247 pages; $5.00

Despite a large amount of research on developing natural language understanding problems little work has been done on evaluating their performance or potential. The evaluations that have been done have been unsystematic and incomplete. This has led to uncertainty and confusion over the accomplishments of natural language processing research.

The lack of evaluation can be primarily attributed to the difficulty of the problem. The desired behavior of natural language processors has not been clearly specified. Partial progress toward the eventual goals for natural language processors has not been delineated much less measured.

This thesis attempts to clarify some of the difficulties behind evaluating the performance of natural language processors. It also proposes an evaluation method that is designed to be systematic and thorough. The method relies on considering a natural language processor from three viewpoints in the light of several taxonomies of issues relevant to natural language processing. Finally an evaluation is described of PLANES a natural language database query system.